

Radiometric alignment and vignetting calibration

Pablo d'Angelo

University of Bielefeld, Technical Faculty, Applied Computer Science
D-33501 Bielefeld, Germany
pablo.dangelo@web.de

Abstract. This paper describes a method to photometrically align registered and overlapping images which have been subject to vignetting (radial light falloff), exposure variations, white balance variation and non-linear camera response. Applications include estimation of vignetting and camera response; vignetting and exposure compensation for image image mosaicing; and creation of high dynamic range mosaics. Compared to previous work white balance changes can be compensated and a computationally efficient algorithm is presented. The method is evaluated with synthetic and real images and is shown to produce better results than comparable methods.

1 Introduction

Grey or colour values captured by most digital and film cameras are related to the scene irradiance through various transformations. The most important ones are the uneven illumination caused by vignetting and a non-linear camera response function. Many applications require the recovery of irradiance, for example irradiance based reconstruction mechanisms, such as Shape from Shading, and Photoconsistency. Image mosaicing is also strongly affected by vignetting. Even if advanced image blending mechanisms are applied [1,2,3], residuals of the vignetting are noticeable as wavy brightness variations, especially in blue sky. Another important use case is the creation of high dynamic range mosaics, where camera response, exposure and vignetting should to be compensated for.

Vignetting is usually corrected by dividing the image by a carefully acquired flat field image. For wide angle or fisheye lenses, this method is not very practical. Additionally knowledge of the camera response function is required for the flatfield correction, which is often unknown if consumer cameras are used. Since vignetting and non-linear camera response lead to spatially varying grey value measurements, they can be calibrated and corrected using grey level pairs extracted from partially overlapping images [1,4]. Litvinov and Schechner [4] estimate non-parametric models of both vignetting and camera response, resulting in a high number of unknowns. Consequently, the algorithm has only been evaluated on image sequences with an overlap of approximately 70-80 percent between consecutive images, requiring many images for the creation of a mosaic. The proposed method is closely related to [1], since it uses similar models for vignetting

and response curve. The vignetting behaviour is modelled by a radial polynomial and camera response function is modelled by the first 5 components of a PCA basis of response functions [5]. Previous work [1,4] on vignetting and exposure correction has assumed equal behaviour of all colour channels, while the proposed approach can correct images with different white balance settings. Goldman and Chen [1] iteratively alternate between vignetting and response estimation, and irradiance estimation. Both require optimisation steps non-linear optimisation, leading to a computationally expensive and slowly converging method. In contrast, the proposed method avoids the direct estimation of scene irradiance by minimising the grey value transfer error between two images, resulting in a minimisation problem with a much lower number of parameters and significantly reduced computation time.

2 Image formation

The image formation used in the paper models how the scene radiance L is related to grey value B measured by a camera. We describe the imaging process using precise radiometric terms. The image irradiance E is proportional to the scene radiance, and given by $E = PL$, where P is a spatially varying attenuation due to vignetting and other effects of the optical system. For simple optical systems, $P = \pi \cos^4 \theta / k^2$, where θ is the angle between the ray of sight and the principle ray from the optical axis and k is the aperture value [6]. The spatially variant $\cos^4 \theta$ term only applies to simple lenses implementing a central projection. For real lenses the spatially variant attenuation often strongly depends on the aperture value k , and needs to be calibrated for each camera setting. We describe the spatially varying attenuation caused by the optical system with a function M .

The irradiance E incident on the image plane is then integrated at each pixel on the imaging sensor with an exposure time t , resulting in an measurement of the energy Q . After scaling the measured energy Q with the a gain factor s , the resulting value is subject to a camera response function f . The grey value measured by the camera is thus given by¹:

$$B = f(eML). \quad (1)$$

For brevity, it is convenient to define an effective exposure $e = \pi st/k^2$, which includes all constant exposure parameters. While scientific cameras usually have a linear response function, most consumer cameras apply a non-linear response function, for example to archive a perceptually uniform encoding. Except for exotic cameras, the camera response f is a monotonically increasing function, and an inverse camera response function $g = f^{-1}$ exists. The radiance of a scene

¹ Real cameras might also suffer from various other systematic effects, for example dark current, as well as spatial and temperature dependent gain variations. These effects are out of the scope of this paper.

point can then be reconstructed from the grey value B by solving Eq. (1) for L :

$$L = \frac{g(B)}{eM}. \quad (2)$$

For colour cameras, we assume that the exposure e of a channel i is scaled by a white balance factor w_i and that the same response function f is applied to each channel.

$$B = f(w_i e M L) \quad (3)$$

To avoid ambiguity between e and w_i , we fix the white balance factor of the green channel to be unity.

2.1 Grey value transfer function

Assume two images of a static scene have been captured with different exposure, or different camera orientation. For the same point, two different grey values B_1 and B_2 are measured due to the different exposures or a spatially varying vignetting term M . We assume that the camera position does not change, and thus the same radiance L is captured by the camera². This results in the following constraint:

$$\frac{g(B_1)}{e_1 M(\mathbf{x}_1)} = \frac{g(B_2)}{e_2 M(\mathbf{x}_2)}. \quad (4)$$

Note that M depends on the positions \mathbf{x}_1 and \mathbf{x}_2 of the corresponding points in the images, and thus cannot be cancelled. By solving equation Eq. (4) for B_1 , we arrive at the grey value transfer function τ :

$$B_1 = \tau(B_2) := f\left(g(B_2) \frac{e_1 M(\mathbf{x}_1)}{e_2 M(\mathbf{x}_2)}\right). \quad (5)$$

By estimating the function τ , it is possible to determine values for the exposures, response curve and vignetting behaviour.

2.2 Parametric models of response curve and vignetting

The camera response function can be represented either by non-parametric [7,8,4] or parametric models [9,5]. In this paper we have used the empirical model of response (EMoR) of Grossmann and Nayar [5], which is a PCA basis created from 201 response curves sampled from different films and cameras. Compared to polynomial and non-parametric models, it contains strong information about the typical shape of camera response curves and suffers less from over fitting and approximation problems. The response curve can then be computed with

$$f = f_0 \sum_l \alpha_l h_l, \quad (6)$$

² If the camera is moved, different radiance values will be captured for objects with non-Lambertian reflectance.

where f_0 is a mean response curve and h_l is the l th principal component of all the response functions considered in the EMoR model. The parameters α_l define the shape of the response curve.

The vignetting M of most lenses can be modelled well with a radial function [1], in this paper we use a radial polynomial:

$$M = \beta_1 r^6 + \beta_2 r^4 + \beta_3 r^2 + 1. \quad (7)$$

Here, $r = \|c - \mathbf{x}\|_2$ is the euclidean distance of a point \mathbf{x} in the image plane from the centre of vignetting c . When using a perfect lenses and camera system, the vignetting centre should coincide with the image centre, however in practice it usually does not, probably due to mounting and assembly tolerances.

3 Estimation of vignetting and exposure from overlapping images

Previous work [4] for the estimation of vignetting, response and exposure are based on directly minimising the error between the two radiance values $g_1(B_1)$ and $g_2(B_2)$ recovered from the grey values B_1 and B_2 measured at a corresponding scene point using a linear method in log space. This formulation has two trivial, physically non-plausible solutions with $e_1 = e_2 = 0$, and $g = \tau = const$. In order to avoid these solutions, soft constraints on shape and smoothness of the response and vignetting behaviour are used. The method presented in [1] minimises $B_1 = f_1(L)$ and $B_2 = f_2(L)$. Since the scene radiance L not known, alternative non-linear estimation steps for of a large number of radiance values L and the vignetting, exposure and response parameters are required, leading to a computationally expensive method.

In this paper we propose to estimate the parameters of the grey value transfer function τ (cf. Eq. (5)) directly. This avoids explicit modelling of the unknown scene radiance L . Grossberg and Nayar [10] use the gray value transfer function to estimate the camera response fuction, but do not consider vignetting. The resulting error term is given by

$$e = d(B_1 - \tau(B_2)), \quad (8)$$

where d is a distance metric, for example the Euclidean norm. Compared to the previously proposed approaches [1,4], our algorithm does not suffer from physically not plausible trivial solutions and only estimates a small number of parameters. The calculated error e is only meaningful if both B_1 and B_2 are well exposed. As shown in [10] and [4], the estimated parameters are subject to an exponential ambiguity, if exposure e and the camera response parameters α are recovered simultaneously.

In the experiments, we have used the Levenberg-Marquardt [11] algorithm to minimise Eq. (8) using least mean squares for all corresponding points between two images. Ordinary least mean squares assumes that only B_1 is subject to Gaussian noise, while B_2 is assumed to be noise free. For the given problem, both

B_1 and B_2 are subject to noise, and an errors-in-variables estimation should be used to obtain an optimal solution. For the results presented in this paper we have used a symmetric error term $d(B_1 - \tau B_2) + d(B_2 - \tau^{-1} B_1)$.

This formulation does not enforce the monotonicity of the camera response. In practice, this is not a problem as long as the measured grey values B_1 and B_2 are not heavily corrupted by outliers. We have found experimentally that for photos with a high proportion of outliers ($> 5\%$), non-monotonous response curves can be estimated by our algorithm. Based on our assumption of monotonously increasing response curves, we require $\frac{\partial f}{\partial x}$ to be positive. Assuming a discrete response curve f defined for grey levels between 0 and 255, a monotonous response curve can be enforced by using

$$e_m = \sum_{i=1}^{255} (\min(f(i) - f(i-1), 0))^2. \quad (9)$$

as an additional constraint in the objective function. While using a penalty function is not the most effective way to enforce this type of constraint, our experiments have shown that Eq. (9) can be effectively used to enforce monotonous response functions without affecting the accuracy of the recovered parameters for scenarios with a reasonable amount of outliers.

Given a set of N corresponding grey values B_1 and B_2 , the parameters e , p , w , and α can be estimated by minimising the error term

$$e = Ne_m + \sum_i^N d(B_{i1} - \tau B_{i2}) + d(B_{i2} - \tau^{-1} B_{i1}) \quad (10)$$

using the Levenberg-Marquardt method[11]. We have explored the use of the Euclidean norm $d(x) = x^2$ and the Huber M-estimator

$$d(x) = \begin{cases} x^2 & |x| < \sigma \\ 2\sigma|x| - \sigma^2 & |x| \geq \sigma \end{cases} \quad (11)$$

4 Application scenarios

The radiometric calibration approach described above is very general since it includes the main parameters influencing the radiance to image grey value transformation parameters. The method has applications in many areas of computer vision, graphics and computational photography.

Many computer and machine vision algorithms, for example 3D shape reconstruction using Shape from Shading or Photoconsistency, expect grey values proportional to the scene radiance L . Even if cameras with linear response function are used, vignetting will lead to an effectively non-linear response across the image. This is especially true if wide angle lenses and large apertures are used, for which it is hard to capture sound flatfield images to correct the vignetting.

The presented method can be used to recover a good approximation of the true vignetting behaviour by analysing as few as 3 overlapping registered images.

Many applications involve the merging of multiple images into a single image. The most prominent example is Image Mosaicing and the creation of panoramic images. In this context, vignetting and exposure differences will lead to grey value mismatches between overlapping images, which lead aesthetically unpleasant results or complicate analysis of the images. In many of these applications, either the response curve of the camera or the exposure of the images are known, allowing an unambiguous determination of the remaining, unknown parameters. For artistic images, the recovery of scene radiance is often not required or even desired. In this case it is often advantageous to reapply the estimated camera response curve after correction of exposure, white balance and vignetting differences, thus sidestepping the exponential ambiguity [1].

5 Experimental results

5.1 Extraction of corresponding points

The estimation of the response and vignetting parameters with Eq. (5) requires the grey value of corresponding points. The panoramic images used for the examples were aligned geometrically using the Hugin software. After registration, corresponding point pairs between the overlapping images are extracted. For real images sequence there will always be misregistration, either due to movement in the scene or camera movement. By choosing corresponding points in areas with low gradients, the number of outliers caused by these small missregistrations can be reduced. Extrapolation problems of the polynomial vignetting term can be avoided by using corresponding points that are roughly uniformly distributed with respect to their radius r .

We sample a set of $5n$ random points and bin them according to their distance from the image centre. The points in each of the 10 bins used are sorted by the sum of the absolute gradient magnitude in the source images. From each bin, only the first $n/10$ points with the lowest gradient values are used for the estimation. This procedure results in corresponding points that are both localised in areas with low image gradients and roughly uniformly distributed with respect to their distance from the image centre.

5.2 Synthetic example

We have analysed the performance of the algorithm using 6 synthetic images that have been extracted from a single panorama, and were transformed into synthetic images with known camera response, exposure and vignetting parameters. After adding Gaussian noise with $\sigma = 2$ grey levels and random outliers, which have been simulated by replacing some simulated grey values with uniformly distributed random numbers. Finally, the synthetic images have been quantised to 8 bit.

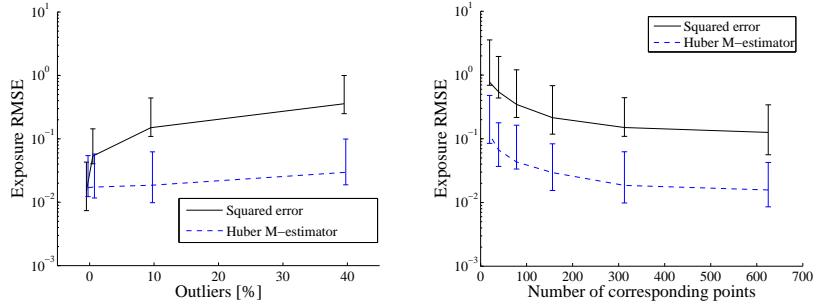


Fig. 1. Evaluation of the robustness on a synthetic 6 image panorama. Left: Exposure error over number of outliers, with 300 corresponding points. Right: Exposure error over number of grey value pairs, with 10% of outliers. The error bars indicate the maximum and minimum deviations against the ground truth measured over 30 simulations.

The proposed algorithm has been used to determine the accuracy of the estimated parameters and to investigate the behaviour of the algorithm with respect to the number of required grey value pairs and the choice of the norm d in Eq. 10. Each experiment was repeated 30 times. Figure 1 illustrates that contrary to [1], we found that using the Huber M-estimator with $\sigma = 5$ grey values results in significantly improved robustness, if the data contains outliers. For larger numbers of outliers, the squared error does not only produce results with larger errors (as expected), but also requires more iterations until convergence. Figure 1 shows the reconstruction error for an increasing amount of grey values pairs. As expected the accuracy of the solution improves if more correspondences are used. If correspondences without outliers and Gaussian noise are used both approaches produce similar results.

5.3 Real examples

We have applied our method to multiple panoramic images sequences, captured with different cameras and under different conditions. Figure 2 shows a panorama consisting of 61 image, captured using a Canon 5D with a manual focus Yashica 300mm lens. The strong vignetting behaviour is corrected almost perfectly. The estimation of vignetting and response curve took approximately 7 seconds, including the time to extract 5000 corresponding points. The panorama shown in Fig. 3 has been captured in aperture priority mode, resulting in images captured with variable exposure time but fixed aperture. The images have been published in [1], and can be used to compare our result to the previous approach. Some very minor seams are still visible, but these can be easily removed by using image blending [2]. Figure 4 shows a $360^\circ \times 180^\circ$ panorama created from 60 images captured with a consumer camera in fully automatic mode. In addition to exposure time changes, the aperture varied between f3.5 and f5.6. The remaining visible

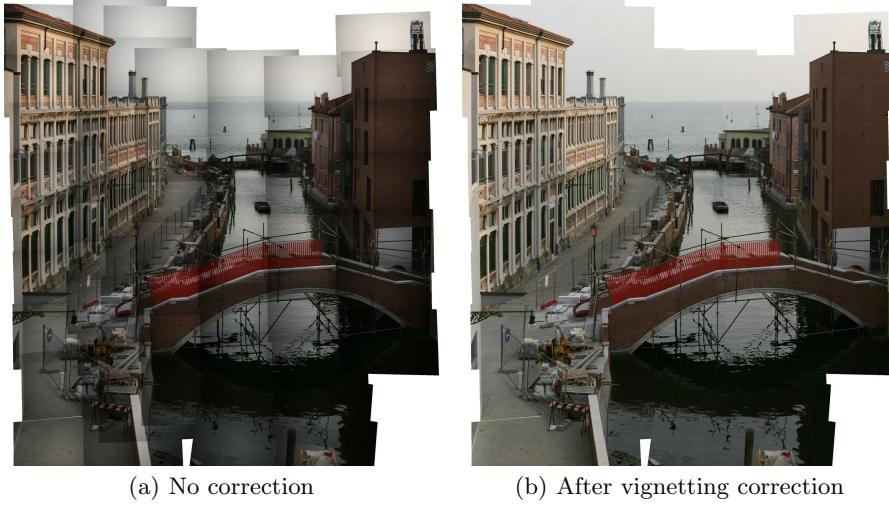


Fig. 2. Venice sequence, 61 images captured with fixed exposure and white balance. As seen in a), the used lens suffers strongly from vignetting. Images courtesy of Jeffrey Martin, <http://www.prague360.com/>

seams in the sky caused by moving clouds. This scene also shows that the proposed algorithm can robustly handle images with some moving objects, as predicted by the synthetic analysis. High resolution images and further evaluation with respect to image blending are available at <http://hugin.sf.net/tech>.

6 Summary and conclusion

We have proposed a method to estimate vignetting, exposure, camera response and white balance from overlapping images. The method has been evaluated on synthetic and real images. It produces accurate results and can be used for vignetting, exposure and colour correction in image mosaicing. If either response or exposure is known, the scene radiance can be recovered. Compared to previous approaches, the method can cope with white balance changes, requires less dense data and has a favourable computational complexity. The described method is implemented in the open source panorama creation software Hugin, available at <http://hugin.sf.net>.

References

1. Goldman, D.B., Chen, J.H.: Vignette and exposure calibration and compensation. In: The 10th IEEE International Conference on Computer Vision. (Oct. 2005) 899–906

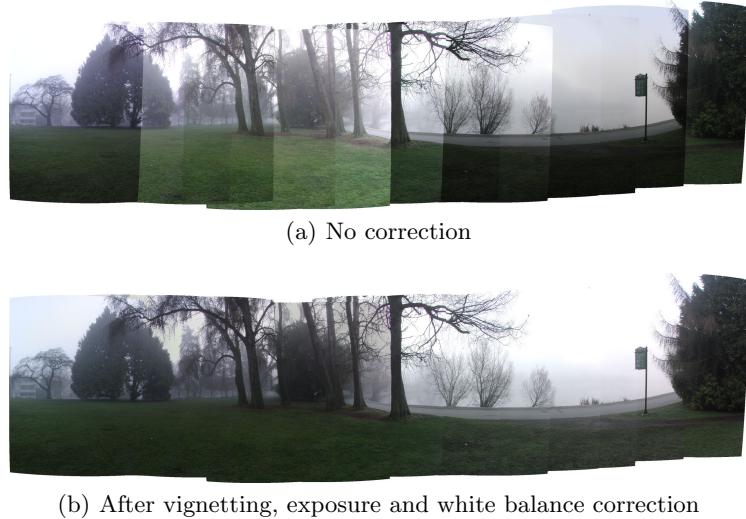


Fig. 3. Foggy lake panorama captured with a consumer camera with unknown response in aperture priority mode, courtesy of [1].

2. Burt, P.J., Adelson, E.H.: A multiresolution spline with application to image mosaics. *ACM Trans. Graph.* **2**(4) (1983) 217–236
3. Agarwala, A., Dontcheva, M., Agrawala, M., Drucker, S., Colburn, A., Curless, B., Salesin, D., Cohen, M.: Interactive digital photomontage. *ACM Trans. Graph.* **23**(3) (2004) 294–302
4. Litvinov, A., Schechner, Y.Y.: Radiometric framework for image mosaicking. *JOSA A* **22**(5) (2005) 839–848
5. Grossberg, M., Nayar, S.: Modeling the Space of Camera Response Functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **26**(10) (Oct 2004) 1272–1282
6. Jähne, B.: *Digital Image Processing*. Springer (2002)
7. Debevec, P.E., Malik, J.: Recovering high dynamic range radiance maps from photographs. In: SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques. (1997) 369–378
8. Robertson, M., Borman, S., Stevenson, R.: Estimation-theoretic approach to dynamic range improvement using multiple exposures. *Journal of Electronic Imaging* **12**(2) (April 2003)
9. Mitsunaga, T., Nayar, S.: Radiometric Self Calibration. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). Volume 1. (Jun 1999) 374–380
10. Grossberg, M., Nayar, S.: Determining the Camera Response from Images: What is Knowable? *IEEE Transactions on Pattern Analysis and Machine Intelligence* **25**(11) (Nov 2003) 1455–1467
11. Lourakis, M.: levmar: Levenberg-marquardt nonlinear least squares algorithms in C/C++. [web page] <http://www.ics.forth.gr/~lourakis/levmar/> (Jul. 2004) [Accessed on 31 Jan. 2007].



(a) No correction



(b) After vignetting, exposure and white balance correction

Fig. 4. Spherical panorama created consisting of 60 images captured with a consumer camera in automatic exposure mode. Some seams are still visible, probably due to different aperture settings. Images courtesy of Alexandre Duret-Lutz, <http://www.flickr.com/photos/gad1/>